



A11106 228923

REFERENCE

NIST
PUBLICATIONSU.S. DEPARTMENT OF COMMERCE
National Institute of Standards and Technology

Information Technology Laboratory

Convergent Information Systems Division

Distributed Systems Technologies Group

Sponsored in part by
Nist's Advanced
Technology
Program

NISTIR 6713 Assurance Hierarchies in B2C Electronic Commerce

G.E. LYON
H.C. TANG

U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Institute of Standards and Technology
Gaithersburg, MD 20899

February 2001

QC
100
.U56
NO. 6713
2001

Assurance Hierarchies in B2C Electronic Commerce

G. E. Lyon

H. C. Tang

U. S. DEPARTMENT OF COMMERCE
Technology Administration
Convergent Information Systems Division
National Institute of Standards
and Technology
Gaithersburg, MD 20899

Sponsored in part by the Advanced
Technology Program, NIST

February 2001



U.S. DEPARTMENT OF COMMERCE
Donald L. Evans, Secretary

NATIONAL INSTITUTE OF STANDARDS
AND TECHNOLOGY
Dr. Karen H. Brown, Acting Director

Assurance Hierarchies in B2C Electronic Commerce*

G. E. Lyon and H. C. Tang

Distributed Systems Technologies Group

Convergent Information Systems Division

National Institute for Standards and Technology

100 Bureau Drive Stop 8951

Gaithersburg MD 20899-8951

{lyon, tang}@nist.gov

Abstract

The fluid, dispersed market of e-commerce makes it hard to judge the reliability of participants. While early Internet trust practices merge old practices—acknowledgments, customer surveys, audits—into Web technology, a need exists for better approaches to assurance. Ideally, these trust mechanisms should be inexpensive, scalable, and private. The proposed three-level hierarchy of ever more powerful assurance services, each interoperable with the other, is one means of addressing the varied assurance requirements of e-commerce participants. Particular emphasis is given to making assurance services inexpensive.

Keywords

ASP; assurance; B2C; customer; electronic commerce; framework; PKI; reputation; seller; World Wide Web.

INTRODUCTION

This document reports on work in progress. As such, it has a dual focus. The first is an attempt to sketch a larger view of assurance—this yields the three-level hierarchy mentioned in the abstract. A second focus employs details and observations arising from a specific project, namely the FAST effort on economical Internet trust services. Salient points from the FAST work punctuate the text to give it a more concrete grounding. This approach may seem discursive, but it is more realistic than just discussing an abstract hierarchy.

DISPERSED MARKET, SMALL CLIENTS

Electronic commerce (*e-commerce*) is defined here as a broad, interdisciplinary field addressing the automation of business via open, globally spanning, Internet public access. E-commerce shows great promise in improving the efficiency of current

business practices and in fostering completely new forms of business transactions. The latter—the ability of the Internet to support novel business mechanisms—has been opened hitherto unrealizable avenues, such as customer-to-customer (C2C) auction markets. These auctions simply could not exist without Internet's powerful aggregating of parties across vast distances and populations. This aggregating is an extremely powerful business force. However, a widely dispersed market brings new challenges in whom to trust. Identity, trust and reputation are essential elements in any business deal. Since the Internet provides no physical presence, business participants must rely upon other means to establish identity and assess reputation.

Initial successes in e-commerce have occurred among large organizations that have resources to build what are often complex and totally closed membership systems. Although this business-to-business (B2B) electronic commerce has had much emphasis, the market for open participation, business-to-customer (B2C) e-commerce is similarly important. Several observations reinforce this view. (1) The larger firms, having greater resources, can initiate projects on far broader scopes than can small firms, so early development might be expected in B2B. (2) The size of B2B is exaggerated. Gross figures, rather than value-added, are commonly tallied for each stage in the B2B distribution chain. (3) Over half of US employment is with establishments classified as small. (4) Smaller businesses have characteristics similar to retail customers, e.g., they are sensitive to cash flow and are unknown to the public. In comparison to regular B2B, the supplies and services that smaller businesses purchase are small—they are close to retail sales in magnitude.

In Web-based B2C neither seller nor customer knows the other, so no basis of trust exists. Distrust persists even when the network is known secure. The true issue for a would-be buyer in this circumstance is whether the *seller is bona fide*. A seller has similar reservations about customers. A Web-based assurance or trust protocol is designed to dispel wariness among

* Report to NIST ATP for project "Assessment of components for micro e-commerce", FY2000.

Mention of commercial products or services is only for illustrative purposes and should not be construed as an endorsement, express or implied.

trading parties who do not know each other. High on the lists of most participants are major concerns about (1) identities of other parties, (2) items being claimed, promised or contracted for in the transaction, and (3) privacy for individuals and businesses. Identity may include establishing the name and true URL for a Web company. The characteristics in question might include a customer's ability to pay and the vendor's record of timely shipping. Regarding privacy, businesses generally are less concerned about transaction anonymity than are private parties, although companies will want to protect proprietary information on markets, price structures, sales and inventories.

The following discussion explores architectures for assurance for the c-commerce process. Particular emphasis is given to scalability and economy of operation, attributes crucial in expanding e-commerce into the domain of very small businesses. This view reflects author participation in a recent industry study on e-commerce identity and attribute authentication. Brief discussions of PC thin clients fit within this focus and are interjected when appropriate. Smaller businesses and some retail customers will want the least expensive, most easily maintained equipment, and this parsimonious view must be accommodated for commercial viability of any assurance method.

ASSURING CLAIMED FACTS

Assurance is not the same as security and it goes beyond establishing identity. In an e-commerce transaction, parties will make various claims about their standing or reputation in the community (*can pay, will ship from inventory*). Whether or not these claims are substantiated—and to what degree—is the task of assurance.

Example I: Audit-Based Assurance

Table 1, below, shows a typical contemporary Web assurance method (see Reference 5 for a functioning site). The mechanism is structured similar to common accounting practices. In step 1, the vendor is audited for, among others, organization of e-commerce systems, accounting practices, general security practices, specific practices on protecting credit card information, and issues of fulfillment (time-to-ship, policy on backorders). When the audit is successfully passed, the third party auditor issues a digital certificate to the seller. This certificate has validity for several months, after which it must be renewed via a checkup audit. The customer sees certification evidence via a displayed seal during an inquiry (step 2) and gets further details (step 3). The customer gets the certificate from the would-be vendor and checks independently (step 4) with the third-party to make sure it is issued validly.

1	Verifier to Seller	Certificate to seller, valid for 3-4 months
2	Customer to Seller	Initial inquiry via Web, sees assurance seal
3	Seller to Customer	Web response (incl. <i>certificate</i>)

4	Customer to Verifier	Checks on certificate details at verifier's site
5	Customer to Seller	Purchase details (includes payment)
6	Seller to Customer	Automatic confirmation

Table 1. Audit-Based Assurance Protocol

Assurance details are checked independently by the customer, outside the merchant's Web site. Assured by reading the certificate, the customer makes an order (step 5) and receives confirmation (step 6). The mechanism is effective only if the customer actually checks certificate details at the remote third party site (step 4). Periodic audits to retain the seal add considerable expense, especially for very small merchants.

Example II: A Trusted Central Service Portal

The Web store portal approaches questions of small seller trust assurance from a completely different angle. For a charge, the portal supplies all necessary e-commerce services and assurance. It is a full-service umbrella, or virtual market, supporting merchants, manufacturers, distributors, and retailers of all sizes [4]. Typically, each participant sells products that ship via some convenient means. Sellers must provide suitable electronic catalogs of their products and services. The umbrella store hosts transactions and customer services, runs a business's Web pages, collects payments and orders shipments. Transaction costs are borne by the portal itself, which charges commissions on sales to pay for the service. Customer assurance comes from dealing with the familiar name of the portal, a "digital intermediary" [6]. Lack of seller name recognition is less a problem, since the participant seller works under the mantle of the portal's reputation and its implied and express guaranties. Buyers at such umbrella stores are typically given a blanket shopping warranty. The assurance underwriter is the portal. A merchant becomes a subordinate affiliate paid commissions, typically 5% to 25% of sales. Profitability for the merchant comes through greater sales volume or manufacturing efficiency—host overheads and commissions are fixed. Despite this drawback, the umbrella store is a highly useful e-commerce architecture.

FRAMEWORKS FOR ASSURANCE

A trust framework in e-commerce must address scalability and cost. A widespread, highly scalable layout will have a hierarchy of levels, each related to the others. A flat, simpler organization may not scale quite so well, but it will cost less originally and in maintenance. Finally, a central service, essentially a clearinghouse, is architecturally minimal and even cheaper. This perspective of different scalable levels and costs will be applied in an examination of assurance and trust mechanisms.

Underlying support architectures in e-commerce vary widely, depending upon objectives, costs and capabilities of host systems. While distributed, peer-to-peer connections (P2P) have gained some notoriety from MP3 music redistribution, by far the

most common e-commerce model is a client/server layout. This is assumed here. Generally, the hard working central server must have a substantial amount of computing resources to sustain the significant traffic demands made upon it. In contrast, the client architecture can be a regular PC, a "thin" PC without much memory or disk, or increasingly, wireless appliances such as cellular telephones and palm-sized devices. Client resources exert an ever-present practical constraint in a discussion of assurance mechanisms. (A sketch of typical "thin" PC client hosts appears in Appendix A.) As mentioned, a simple, economical (client) system is very attractive to the small merchant or at-home shopper.

Much like the rest of the Internet, transaction assurance is today in a state of "cut and try." Three assurance organizational forms will be discussed here. (1) The tree-structured public key infrastructure (PKI) and its commercial variants typify an hierarchical architecture that scales up well. (2) A flat non-hierarchical arrangement can be simple and economical. (3) A central, service portal supplies a set of inexpensive, tailored services to member customers; it also serves as a gateway to other services (say from 2 or 1). The hierarchical approach (item 1, above) is being deployed, although costs currently limit it to medium and large companies [1]. The flat assurance structure is the result of a recent study by banking industry associates to make assurance services affordable to smaller participants. Implementation is in planning stages [2]. The combination of PKI's tree hierarchy and a flat, cheaper structure has been discussed as a method of trading of scalability and cost. To further supplement this, a "local" portal is proposed here as an even cheaper method of accommodating the smallest of firms and clients. This gives three levels of assurance mechanisms. A specialized local portal (cheapest) might connect to a flat assurance network (intermediate scalability), which, in turn, links into a broader, more costly PKI-based system. In all cases, interoperability among the various assurance levels is necessary to assure a truly global reach.

Tree

One mature methodology for identity authentication is the Public Key Infrastructure (PKI)[1, 3]. A great strength of PKI is its scalability, which arises from its tree layout. PKI has a hierarchy of authorities that issue digital certificates. The certificate is a claim of identity that also states who issued it. When parties want to check the identity of another (say X), they follow a backward chain from X's certificate until they reach an issuer they share in common with X. (The issuer of X's certificate will themselves have an identity certificate from some source.) In practice, PKI has had problems. For instance, two Internet parties may have certificates issued from completely distinct root certificate authorities (CAs)—there is no agreed world authority or master root. Also, different authorities may enroll members differently. In the case above with separate CAs, the two (or more) distinct issuers must establish a link at some level reachable by all concerned parties. This cross certification is slow relative to the pace of Web events.

However well the tree-structured organization of PKI works, it suffers problems of economics (it is expensive) and limits in scope—certificate identity may be only the first part of what is

needed to pursue a transaction when the body of participants is very large. An unknown entity, even if certified correct for identity, may still trigger no recognition to someone on the Internet thousands of miles away. In a world context, it is highly likely for participants to have different root certification authorities (CAs). As mentioned, there may be no quick, convenient way to reconcile these CAs.

The Flat Framework of FAST

Looking for simplified approaches, the Financial Services Technology Consortium has recently studied an essentially flat framework called FAST—the Financial Agent Secure Transaction [2]. In this framework, users link directly to agent institutions that are members of a closed trust system. The form of the link from customer to agent institution may vary from institution to institution, depending upon circumstances and market. In each case, the institution serves as a fiduciary and assumes some agreed upon level of risk.

Rather than tracing back hierarchical lines of authority as in PKI, FAST has a much simplified, dual network, one part being the usual Internet and a second link being private among the agents of trust [2]. (See Figure 1, below.) Each entity (client, customer, merchant, business, vendor, etc.) enters into Web business transactions in association with their agent of trust. Regular elements of business are conducted as usual over the open Internet. However, all trust agents are tied together by a closed network; agents trust each other to a degree stipulated in a network agreement. There is risk in assurance and the agreement must spell out who assumes responsibility for any loss.¹ A typical trust agent would be a bank, credit union, or, in some countries, a post office. Because FAST is quite interesting and not well known, more details of it will be described in the discussion than for PKI or a portal.

Imagine two FAST users, client 1 and client 2 of Figure 1, who are attempting to strike a business deal. Internet dealing then involves two links for each. The upper link is the usual Internet dialogue that carries everyday business negotiations. The bulk of the negotiation will occur here. Two secondary links connect clients 1 and 2 to their respective trust agents, who are themselves linked in a closed-membership trust network arrangement. Critical items, such as client 2's true URL, can be checked by client 1 asking through agent 1 into the trust back channel. In contrast to PKI architecture, which elaborates hierarchically, the FAST trust structure is essentially flat; it relies upon its single level of trust agents to handle the far more numerous public clients (businesses and their customers). The FAST trust network would piggyback upon an existing, well-tested system, such as the ACH (Automated Clearing House) network that banks and credit unions use (other networks could be employed [10]).

¹ E.g., North American merchants assume responsibility for identity verification in credit card purchases. An improper identification allowing fraud becomes a loss to the merchant.




Figure 1. Trust Questioning via a Closed Network

Business entities (merchant or customer) link to their trust agents in any manner the agents deem appropriate for the transaction and risk. Assurance responsibility will lie with each agent. This risk can be offset by combinations of fee levels and levels of security to the client. Transactions fees reflect loss possibilities and will be determined by each trust agent. No compatibility in trust mechanism linkages need exist between two clients. Only agents exchange trust information, using FAST. The trust network secures agents together and guarantees that what client 2 hears from agent 2 is what agent 1 said about client 1. Notice that this context underscores the importance in e-commerce architecture of clearly determining procedural rules, regulations and responsibilities. These rules governing operations can be as important as any implementation details.

Single Portal

Although FAST was developed in response to the expense of PKI, there is an opportunity for an even simpler approach for certain classes of naïve commercial users. The umbrella e-store mentioned earlier serves as illustration. A central service portal guarantees elements of an e-commerce transaction. Here, the user relies upon knowing the portal's reputation and trusting it, even though individuals or small businesses using the site are completely unknown. Several commercial versions of this framework have been tried. Reference 4 points to an example. Imagine the FAST agents 1 and 2 (all FIs and other trust agents) collapsed into one central point of service. This is the structure of the Custom Service Portal (CSP)--it serves *all* clients. A very small, local manufacturer might be such a client. The advantages of a CSP are primarily for local users, who while unfamiliar

with each other's reputations, place significant trust in mediation supplied by the portal, which is responsible for the honest and reliable enrollment of all the clients. Such a portal ties together manufacturers, material suppliers and a host of related participants. Service demands upon the portal might range from light—a manufacturer posts its catalog on the server and runs Web-based sales—to substantial. In the latter, the user might want opportunities for making background checks on potential suppliers, handling biddings for business, exchanging materials, and even managing contractual elements. The CSP is essentially a application service provider (ASP) additionally tailored for assurance, identity and trust functions for a community.

FAST is seen (initially) as a network for the financial industry, an inexpensive service built *ad hoc* upon the available, secure ACH (automated clearing house) network. In contrast, each CSP would be tailored to its unique customer base. A CSP might actually be sponsored by a regional governmental authority as an inexpensive cooperative that fosters market growth and promotes good manufacturing and marketing practices. Given this mixed sponsorship (CSPs might vary widely across geographic regions of the US), it is important provision be made that all such CSPs can talk to each other. The FAST is designed for this. FAST, servicing smaller business participants, would in turn link into PKI systems to achieve a more global reach in e-commerce. This would give an interoperable global network for assurance. See Figure 2, immediately below.




Figure 2. A Hierarchy of Assurance

A FAST PERSPECTIVE ON RELATED CONCERNS

Privacy is always a concern of Web-based users. When payments are involved, protection of credit card numbers is certainly needed. In Europe, where debit cards are far more common, the need for economic privacy is even higher, since these cards are less restricted in the risk they pose to the user. Privacy can have dimensions that are technical (security), policy (trust) or economic (cost and profit) [7]. There is always a tradeoff among these. Even excellent security/privacy protocols have to survive a business case analysis. Full public key methods (e.g., the PKI structures mentioned earlier) are probably too complex for most small Web retailers. Other

approaches for secure and private electronic transactions have also proven too expensive and failed the business case: U.S. merchants currently find it cheaper to absorb losses incurred with lesser levels of protection than to use more secure but more expensive protocols [8, 3].

Privacy

Figure 3 (below) illustrates FAST's privacy feature—it depicts the framework with one set of questions that a seller might need resolved for the sale of a case of wine (example from the FAST Phase 1 draft report [2]). The seller (upper right corner) is worried about minimal age for the purchaser's state and whether payment is likely. The buyer too will have questions—whether the URL of the site is the correct one and whether reliability of the seller is satisfactory (for clarity, these are omitted from the figure). Two copies of the seller's questions are sent, one to Client 1 for approval and the second through Agent 2 to Agent 1 for comparison with the approved copy from Client 1. Notice that no archival data on the buyer's age or bank balance is sent—only an answer "yes" as the answer to the conjunction of both questions. The answer is routed through FAST to Agent 2 and then to Client 2—the seller. Client 1 in Figure 3 is identified only by an identification number (#2998) generated for the transaction. Each new deal will generate a new identification number. This makes snooping and correlating of Web traffic far less productive. (Some software is also available for privacy on the client-to-client Internet link.) A random ID is similarly available to Client 2 unless the merchant chooses to be known, which most will.

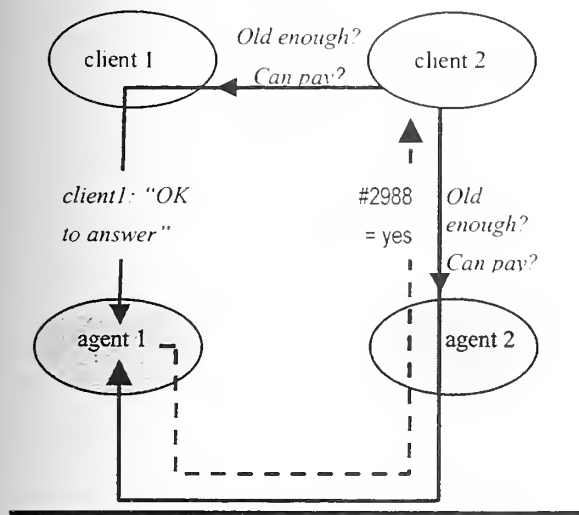


Figure 3. Handling of Inquiries

Payment Methods

The question of methods of payment arose early in FAST. It is easy to argue the business case for electronic payment systems. Observe that with a bank teller each transaction averages \$1.25, whereas each swipe of a credit card

costs about \$0.01.² This two decimal orders of magnitude advantage continues to spur interest in e-payments. Over a dozen new methods of electronic payment have been proposed.³ In addition, a NIST/INT study of an original e-payment scheme called *certified transfer* underscored the importance of keeping FAST independent of any particular payment method [13]. This highlights what is a general problem in e-commerce. The attractiveness of specific system components and configurations changes rapidly. It is crucial that an e-commerce architecture provide for substitution of components, that rapid system reconfiguration be possible. This ideal state of modularity has common names and meanings of data fields, interoperable protocols, and systems components for which inter-component filters are unnecessary. Of course, specific technical implications and constraints may give these laudable objectives a less than ideal realization.

Implementation

FAST as a protocol would include the application, presentation and session layers of the conventional communication protocol stack. The FAST *message protocol*—customer to merchant, customer/merchant to financial institution (FI) and financial institution to service provider, are in the application and presentation layers. The FAST *service protocol* would occupy presentation and session layers and use sockets, TCP/IP and Internet. The service protocol would be most portable. The FAST message protocol might depend upon sockets, TCP/IP, Internet as well as a proprietary communications infrastructure. However, the proprietary links would be isolated within customer services of a particular FI—it would be where the FI differentiated its service from that of its competitors. This differentiation would be isolated and not interfere with FAST interoperability.

Client Architectures

Client architectures were left open in FAST discussions. However, a couple of persistent constraints did arise from deliberations. First, FAST client-side (in contrast to agent/server side) software should be kept to a minimum. Clients dislike complex software on their systems—it promotes systems errors and pushes costs up. Much FAST discussion centered around having just a browser on client systems. This has some weaknesses, however. Certain "man in the middle" attacks on the open Internet link cannot be detected unless there is software to record SSL session numbers: Remember that the FAST model supports clients who are *completely unknown* to each other—a digital certificate might not signify much (conventional "man-in-the-middle" attacks are caught by the sender's certificate being wrong).

² D. Schutzer, presentation, "Internet Risk Management", FSTC Fall General Meeting, 2 October 2000, Bethesda MD.

³ Including B2B, there are nearly forty payment solutions offered for e-commerce. Most are computerized versions of earlier financial practices and many suffer deficiencies that will keep them from general use. Fourteen or so are second-generation Internet payment instruments (T.P. Vartanian, presentation, FSTC Fall meeting, October 2000).

A second assumption was that in two years or sooner, cellular wireless clients would be heavy users of such a FAST-like framework.⁴ Cellular appliances are certainly limited in their resources. The usual browser now becomes an even weaker mini-browser. Intelligent agents in proxy servers can help, but this opens security problems. There is already a problem with proxies, which appear because a user's wireless link is only local. Within each locality, the user's broadcast signal is received and converted by a proxy into a transaction for the wired Internet. The problem arises when these wireless users want information behind firewalls. An example would be a call to the home office for latest figures in some engineering project bid being developed by a traveling team scattered over the USA. One does not want the information in some shadow copy outside the firewall—it is highly vulnerable and not necessarily up to date. Concern now lies with how to let the foreign entity—the wireless carrier proxy acting in a (supposedly) authorized fashion—through a firewall. A last observation is that wireless appliances may generate a strong need for biometrics, since loss of such a small, powerful device might otherwise jeopardize critical elements in an organization.

The client/server model thrives because it handles many problems that customers want solved. ASP (application service providers) understand this and have been vigorously positioning their offerings to attract clients. For many valid reasons, small merchants and consumers are not fond of supporting software on their machines. The ASP market has attracted premier computer corporations. Nonetheless, one cannot discount peer-to-peer (P2P) e-commerce models. Because P2P models remove or disperse server functions, they demand a network of more capable (fatter) clients to provide in aggregate what was once a server service. However, as Appendix A shows, competition in PCs has driven prices into very attractive ranges. Hardware cost is not likely an immediate issue for a conventional client system. While the viability of MP3 client-based distribution on the Web may be determined by courts, other client-based services may evolve. Some distributed systems have no central service of any kind. Where to start poses the user a challenge: The whole system is dispersed onto client machines, without any server intervention. Exchanged files can be of any type, not just MP3. Client machines for these software systems must have adequate storage, communications and processing capacities.

DISCUSSION

Having examined briefly a hierarchy of assurance for small e-commerce, a number of observations can be made.

PK Architecture

⁴ Several brokerage firms do allow trading via wireless Internet access. Concerns exist, however, about snooping—wireless is much easier to monitor. For now, the companies accept the higher risk. This acceptance of risk is true for e-banking as well; accounts have disappeared and they have been the banks' loss.

Public Key (PK) methods will continue to advance into the market, even though the pace at present has not been nearly as swift as some had predicted or hoped. Deployment has been halting in many banks, where technology officers who purchased large numbers of digital certificates have found organizational demand for certificates disappointingly low. However, advancing technologies will make public keys and certificates more popular. Take for example biometric identifiers, which are extremely attractive to both merchant and customer. Such an identifier is unique, with nothing to memorize or forget. The problem arises in making the biometric recognition template for a person's finger, iris, or voice a conventional shared-secret key. What happens if this secret key is stolen? Many security breaches are inside leaks, so it is no idle question. While conventional alphanumeric secret passwords can be changed, this is not an option for a person's biometric identity. The solution is to make the biometric template the *public* key in a person's digital certificate. The person has a corresponding private key to establish identity and the private key can be changed if a security breach occurs. Whether the full PKI mechanism will be employed everywhere is debatable. Pressures of economy may force more simplified versions that, *e.g.*, trade some scalability for ease of maintenance or economy.

Privacy Mechanisms

One of the lessons of FAST is the attractiveness of its security mechanism. A broad list of things might need authentication—citizenship, marital status, education, membership, licensing standings. Checks on legal matters are important in some lines of business—indictments, lawsuits brought by and against the entity in question, criminal records, civil judgments and related settlements. Some of these topics pose delicate questions of privacy for the respondents. Although the FAST system is optimized for cheap, small-scale e-commerce assurance, its central feature for information access provides fewer likely places for disclosure of sensitive materials. No original archival data flows through the FAST network and use of the material can be easily logged. Questions can be aggregated and sent to the answering agent for a "yes" or "no", so removed from context, the actual answer may signify little. Even here, the answer can be attached to the transaction identifier of an anonymous participant. This has a flavor of some of the better security solutions where users are not given secret passwords, but instead, have individual encryption algorithms that reside on their machines. The server sends a random number to the client system, which returns the number encrypted. The actual thing of interest to any thief, the encryption algorithm (in FAST, the archival data), never enters the network. The dual networks in FAST, one open and the other closed, also make the overall framework less accessible to recreational hackers.

Inexpensive Customer Services

The form of e-commerce service for an inexperienced vendor or small manufacturer may require a significant amount of custom tailoring and training. Providing this may prove unattractive for conventional ASP providers; the limited prospects for broader sales of a specialized service may make handling such customers uneconomic. Local groups of merchant may find that forming

modest joint ventures to build custom service portals (CSPs) attractive. Government—local, state and even national—may want to assist such ventures. As expected in business, costs and law influence e-commerce architectures. An overemphasis upon security to the detriment of economy of operation has already limited one approach in secure electronic transfer in the United States. (See Reference 3, p.173-175 for discussion of such a protocol.) The interaction of technical capabilities, business models and customer acceptance is a subtle, not yet mature element in e-commerce.

CONCLUSIONS

Assurance frameworks are evolving rapidly because a need exists for them. Assurance is not the same as security, although it uses security mechanisms as given components. While early trust practices merged old practices (acknowledgments, customer surveys, and audits) into Web technology, a need exists for more sophisticated approaches. Ideally, one wants a framework that is inexpensive, scalable, private and modular. A hierarchy of ever more powerful services, each interoperable with the other, provides one means of addressing these demanding requirements. While a two level structure of PKI and FAST has been discussed between these respective communities, the proposal here is to extend the structure to three levels. The lowest level, serviced by trusted portals, links into something like FAST when participants are not members of the same portal service.

ACKNOWLEDGMENTS

We thank Dr. Len Gallagher for his careful reading and well-taken suggestions on revising the draft version.

REFERENCES

- [1] Identrus. Stresses public key infrastructure compatibility for commercial entities. <http://www.identrus.com>
- [2] Financial Agent Secure Transaction (FAST) Phase I Final Report White Paper, Sept. 2000, 184pp. Financial Services Technology Consortium. <http://www.fstc.org>
- [3] Ford, W. and Baum, M.S. Secure Electronic Commerce, Prentice Hall, Upper Saddle River NJ, 1997.
- [4] Vstore.com organization. <http://www.vstore.com>
- [5] WebTrust™ description text. <http://www.aicpa.org/webtrust/into.htm>
- [6] Ba, S., Whinston, A. B. and Zhang, H. Small business in the digital economy: Digital company of the future. Proceedings of the Conference on Understanding the Digital Economy: Data, Tools and Research (Washington, D.C. 1999), MIT Press.
- [7] Ghosh, A.K. E-Commerce Security. John Wiley & Sons, Inc., New York NY, 1998.
- [8] Secure Electronic Transaction (SET). Reference 3 above has clear information on SET. On the Web, one might look at <http://whatis.techtarget.com> under their Internet heading.
- [9] www.comercenet.com/research/reports/1998/98_13_b.htm
- [10] The Swift network provides secure messaging for 3000 financial members. www.swift.com

[11] Lyon, G.E. Assurance protocols and small Web retailers, Proceedings of the 2000 ACM Symposium on Applied Computing, Como, Italy, March 19-21, 2000, pp. 904-908.

[12] P3P is a privacy protocol effort

<http://www.w3.org/P3P/brochure.html>

[13] M. Meniszez, "Some Evolving Aspects for American E-Commerce", Study Report for the Third Program Year (Institut National des Telecommunications, Evry, France, January, 2000) 88pp.

APPENDIX A. CLIENT ARCHITECTURES

This attachment reviews some client architecture characteristics for small-vendor e-commerce. The text identifies and evaluates these components, gives some illustrations and provides a brief assessment. Generally, the smaller merchants and retail customers will want a service that can run on such machines, even if not all transactions are conducted from them.

Using the open World Wide Web, E-commerce connects a consumer with a retailer, or a supplier with a manufacturer. The simplest servers provide information about products and services. The client surfs the Web to query and retrieve the needed information. More complex sites allow the customer to make purchases or engage in other contractual agreements. Users thus conduct business transactions over the Web. Making a business transaction work requires a number of hardware and software components. A successful e-commerce business also requires more stringent components to support robustness and security.

Client architectures can be "fat" or "thin". A fat client host is usually a full computer with all its flexibilities. "Thin" clients are more restricted; they are more than plain terminals, but their computing or storage is limited. Client architecture used in or available for small vendor e-commerce can be subdivided into three sub-categories: hardware, software, and stringent components. Hardware will vary—thin client to fat client—from a simple, modest network (diskless) computer to a sophisticated computer system. The corresponding software ranges from just a browser for a simple consumer process to a complicated transaction system for business-to-business applications with complex interactions. Stringent components run from simple user-authentication to complicated network security systems, such as implementations of a firewall and SSL protocols. The following sections discuss each of the component dimensions.

Hardware

A client system requires only modest CPU power, memory and disk storage. The hard-drive on the client is for caching accessed web documents, locally edited texts and similar transient data. Server sites run application programs and store more archival data, which eases greatly demands upon the client system. Available PC processors on the market have more capability than required for a thin client. A 200 MHz CPU would provide more than sufficient CPU power for a thin client; current PCs commonly have CPU speeds ranging from 300 to 800 MHz. A thin client probably needs no more than 32M of memory. But

the cost of memory has become a small factor. The performance improvement for a thin client with 64Mbytes of memory would outweigh the small incremental cost of the additional 32M bytes. However, the client site does require high-speed communication bandwidth and graphics display to retrieve and display effectively a lengthy document with images. A thin client should have at least the prevailing 56Kbit/s modem or 10/100 base-T Ethernet and a high-resolution graphic display. The display is usually color.

Software

The client's operating system (OS) can be modest. A Web browser is the primary required software, although other additional software adds benefits. E-mail software is highly desirable if it is not already included in the browser. Text editing software should also be included. For a general on-line shopping system, shopping cart systems, payment systems, inventory database, inventory management system, and accounting system are all necessary, but these can be kept on and run on the server. What complexity the client software should have depends on user's need and budget. For example, a "kidnapster"-like distributed e-commerce model tends to remove (or disperse) server functions, but it definitely increases needed client support, especially for disk storage.

Stringent Components

Security is the most critical issue in e-commerce. Having a strong security system means having a chance for a trusted e-commerce business. A responsive, reliable, secure transaction is a must; thus, the components of the client architecture characteristics must achieve these goals. The Secure Socket Layer (SSL) protocol is included in most commercial Web servers and is supported by most browsers. SSL includes authentication of the connection via certificates, privacy using encryption, and message integrity by message digest. The HTTP/SSL may provide secure transaction between client and server.

Other than the security protocol, there are access control techniques: IPX-to-IP gateways, firewalls, proxy servers, virtual private networks, and tunneling to keep servers secure from unauthorized access. Multimedia requires not only high-speed computational capacity but also great amounts of memory. The thin client may fail either of these requirements for some applications. Whether a thin client should support the requirements of full graphics, audio, and video capabilities is a tradeoff issue. Client resource tradeoffs include cost, reliable hardware and software, session functionality and acceptability, portability, power conservation, and connectivity bandwidth. The application will dictate the balance among these requirements. Clearly, a thin-client platform will run into performance problems when running features of full-blown PCs, such as playing AVI and other video files. It is a tradeoff between need, cost and capability.

Wireless Connections

Wireless connections to the Internet already exist with limited feature hand-held devices and will soon be widely available on thin clients (e.g., small laptops) as well. Mini-browsers, though they may be in early deployment stages, exist in cell phones,

PDAs, pagers, and other wireless devices. The Wireless Application Protocol (WAP) standard at this time is still considered vulnerable to attack at the WAP gateway server. There is a small threat that a sophisticated hacker could enter a WAP gateway server to steal data. Despite this possibility, the level of security is nevertheless considered quite high. (Actual WAP use is still very low.) WML, an XML dialect language for wireless document markup, is being developed and promulgated. J2ME (Java 2 Micro Edition) enables dynamic personalized interactive services for wireless devices.

Examples of Existing Thin Clients

The following are two descriptions of thin client products selected from Internet listings (August, 2000).

500 MHz System -- \$299	
<u>Hardware</u>	
AMD K6-2 500 MHz MMX CPU with 3D Now (OEM)(100MHz bus)	
64 Mbyte SDRAM Memory	
512K 64bit L2 Cache	
56K FaxModem	
Network Adaptor: Fast Ethernet 10/100 Base-T LAN onboard	
16bit PC/Sound Pro 3D Surround Sound with Stereo Mixer	
Integrated 64bit 8 MB AGP 3D Video	
64bit 3D VGA Graphics Accelerator	
Monitor and Speakers are not included.	
2 USB Ports -- PS/2 Port -- 1 serial Port -- 1 ECP/EPP Port	
Hardware Manager detects CPU Temp. and Speed	
Mini-Mid A Tower Deluxe Case with w 250 watt power supply	
<u>Software</u>	
Windows 98 or Linux	
WordPerfect8: Includes Corel Wordperfect 8, Quattro Pro 8, Presentation 8, Corel CENTRAL 8	
GAMUT 98 A professional audio station (exploits the latest MP3 technology)	
PC-Cillin 98 Virus Protection (OEM)	

366 MHz System -- \$127
<u>Hardware</u>
Monitor and Speakers are not included
366MHz Intel Celeron Processor (128K Cache)
Memory: 32MB SDRAM (4mb of which used for video)
Hard Drive: 4.3 G EIDE
1.44MB Floppy Drive
CD ROM(32x)
Featuring a modem master 9050 56K v.90 Modem
Trident 4D Wave PCI Sound
SIS 3D 4MB AG(UMA) video
<u>Software</u>
Windows 98
Manual/license and backup software not included for Windows 98

Observations

The effects of the thin client resource tradeoffs depend upon user requirements—a user with sophisticated operations will certainly need extended functions. The following table lists the range of hardware components likely to be used for an ordinary thin client.

CPU power	200-400 MHz
Memory	32-128 MBytes
Hard drive storage	2-4 GBytes
Network connection	56K Modem or 10/100 Ethernet
Operating System	Windows CE or Equivalent OS
Browser	Netscape or Internet Explorer
E-mail	Netscape/MS Outlook/Eudora

